

Computing Personalised Treatments through *In Silico* Clinical Trials.

A Case Study on Downregulation in Assisted Reproduction.

T. Mancini¹, F. Mari¹, A. Massini¹, I. Melatti¹, I. Salvo¹, S. Sinisi¹,
E. Tronci¹, R. Ehrig², S. Röblitz², and B. Leeners³

¹ Computer Science Department, Sapienza University of Rome

² Computational Systems Biology Group, Zuse Institute Berlin

³ Division of Reproductive Endocrinology, University Hospital Zurich

Abstract. In-Silico Clinical Trials (ISCT), *i.e.*, clinical experimental campaigns carried out by means of computer simulations, hold the promise to decrease time and cost for the safety and efficacy assessment of pharmacological treatments, reduce the need for animal and human testing, and enable precision medicine.

In this paper we present a case study aiming at quantifying, by means of a multi-arm ISCT supervised by intelligent search, the potential impact of precision medicine approaches on a real pharmacological treatment, namely the downregulation phase of a complex clinical protocol for assisted reproduction.

1 Introduction

Model-based approaches to safety and efficacy assessment of pharmacological treatments (In-Silico Clinical Trials, ISCT) hold the promise to *decrease* time and cost for the needed experimentations, *reduce* the need for animal and human testing, and enable *precision* medicine, where *personalised* treatments *optimised* for each patient can be designed before being actually administered. This is to be achieved by developing *computational models* for human physiology, pathophysiology, Pharmacokinetics (PK) and Pharmacodynamics (PD), which also define the possible physiological differences between different individuals (*i.e.*, the possible *phenotypes*).

Research in Virtual Physiological Human (VPH) provides mechanistic quantitative models of the human physiology at levels of scale ranging from body compartments (see, *e.g.*, [8,36]), organs (see, *e.g.*, [7]) down to cells (see, *e.g.*, [2]) or even molecules (see, *e.g.*, [37]).

Independently of the level of scale used, most VPH models are defined by means of complex, highly non-linear differential equations. By means of numerical integration, such models can be executed within *simulators* in order to provide quantitative information about the time evolution of the modelled biological quantities (*e.g.*, hormone concentrations).

Also, VPH models of different body compartments can be integrated in larger models, eventually defining *whole human models*. Important attempts in this direction are given by HumMod [17] and the more recent Physiome [33], a complex whole-body model written in the standard Modelica language and executable by means of open-source (*e.g.*, OpenModelica, openmodelica.org) as well as proprietary (*e.g.*, Dymola, dymola.com) simulators.

1.1 Motivations

Current pharmacological treatments are often designed with the *average* patient in mind. A key topic in *precision medicine* is to develop pharmacological treatments *optimised* for any given individual (*personalised* treatments). Several optimisation criteria can be defined. A typical criterion is minimisation of the overall amount of administered drug, which often also reduces the probability of adverse side-effects and contributes to decreasing the treatment cost.

With their amenability to define different phenotypes, VPH models of proved accuracy are a *key enabler* for precision medicine.

1.2 Contributions

In this paper we present a case study aiming at quantifying, by means of extensive computer simulation-based experimental campaigns (In-Silico Clinical Trials, ISCT) guided by intelligent search, the impact of precision medicine approaches on a real pharmacological treatment, namely the *downregulation* phase of a complex clinical protocol for assisted reproduction in humans. To do this, we exploit a large VPH model of the Hypothalamic-Pituitary-Gonadal (HPG) axis [36] in order to conduct a multi-arm ISCT. This model has been used as a case study in [39,31], where a parallel algorithm based on statistical model checking techniques and deployed on a High Performance Computing (HPC) infrastructure has been described, which computes a representative set of Virtual Phenotypes (VPs) for the VPH model given as input.

Here we select, from the population of VPs computed in [39,31] for the HPG axis model [36], a representative subset of 98 different VPs of our HPG axis model, and conduct a *distinct arm* of our ISCT *for each* of such phenotypes, in order to compute the *lightest* (in terms of overall amount of administered drug) treatment still *effective* for *that* phenotype.

Our multi-arm ISCT has been conducted on a large HPC infrastructure using a backtracking-based search algorithm to seek for the lightest but still effective treatment for each VP. Our search algorithm drives a (black-box) *simulator* of the VPH model at hand to intelligently explore the space of possible time-series of drug administrations.

We note that a clinical trial with a so high number of arms is *not even conceivable* in the classical *in vivo* setting (it would require a distinct human patient for each phenotype *and* for each possible alternative time-series of drug administrations). This undoubtedly shows the *revolutionary* potential of artificial intelligence for model-based (*in silico*) precision medicine.

Paper outline. The paper is organised as follows. After describing some background knowledge on the HPG axis model used in our case study and its phenotypes in [Section 2](#), we outline our reference downregulation protocol in [Section 3](#). In [Section 4](#) we state our problem and describe our algorithm to solve it. [Section 5](#) describes and analyses the outcome of our multi-arm ISCT. Finally, [Section 6](#) discusses related work and [Section 7](#) draws conclusions.

2 Preliminaries

In this section we give the necessary background knowledge.

2.1 The GynCycle Virtual Physiological Human Model

We focus on the GynCycle Virtual Physiological Human (VPH) model, initially presented in [36]. GynCycle is a continuous-time differential-algebraic equation-based mathematical model of the human female Hypothalamic–Pituitary–Gonadal (HPG) axis, with a special focus on the interactions and feedback mechanisms between hormones like GnRH, Follicle-Stimulating Hormone (FSH), Luteinizing Hormone (LH), Estradiol (E2), Progesterone (P4), Inhibin A (IhA) and Inhibin B (IhB) during the different stages of the human female menstrual cycle.

The model defines the time evolution of overall 33 biological quantities (mostly blood concentration of hormones) and the Pharmacokinetics/Pharmacodynamics (PKPD) of several pharmaceutical drugs used in assisted reproduction (in particular, GnRH analogues like Triptorelin, or FSH/LH-based drugs) by means of highly non-linear differential equations.

As such, GynCycle is a *hybrid system* whose inputs represent drug administrations. As it happens with VPH models of practical interest as GynCycle, the complexity of the differential equations makes their symbolic analysis (by, *e.g.*, a model checker for hybrid systems) very challenging. Indeed, the model can be *simulated* by means of numerical integration in order to compute the time evolution of the hormones of interest upon administration of a sequence of doses of one or more drugs. Thus, GynCycle is used as an *executable* black-box model.

2.2 GynCycle Virtual Phenotypes

VPH models like GynCycle typically take into account inter-subject variability (*i.e.*, the physiological differences among different individuals) by including suitable *parameters* in their equations. Different value assignments to model parameters yield different model time evolutions and/or different reactions to drug administrations, thus defining different Virtual Phenotypes (VPs). Intuitively, each VP represents a class of patients behaving similarly to each other.

Computing a complete set of VPs for the model at hand is the starting point to obtain a *representative* population of virtual patients (hence, ideally showing all possible phenotypes). In turn, this is the key enabler to perform In-Silico

Clinical Trials (ISCT) to assess, *e.g.*, safety and/or efficacy of a pharmacological treatment.

Unfortunately, computing the complete set of VPs defined by a complex VPH model is all but easy from the combinatorial point of view, and may require weeks of massive simulation on large High Performance Computing (HPC) infrastructures. In [39,31] a statistical model checking-based approach using a non-uniform sampling process finely tuned against the model has been discussed, and a representative population of VPs for GynCycle consisting of around 10^4 individuals has been computed (*once and for all*) from a huge search space of 10^{75} possible candidates.

The availability of such a representative population of VPs is the *key enabler* for our ISCT.

3 Downregulation in Assisted Reproduction Treatments

In this section we define the treatment we used as our case study: the downregulation phase of an assisted reproduction protocol currently administered at the Department of Reproductive Endocrinology of University Hospital Zurich.

At each menstrual cycle of a human female, among the several follicles initially present in the ovaries, only one, unless in exceptional cases, grows and reaches maturity. A complex competition among follicles takes place, driven by a hormone-based feedback loop, in order to inhibit the maturation of multiple follicles.

In a so-called “long protocol”, the treatment aims at neutralising (through drugs) such a feedback loop (*downregulation phase*), in order to achieve a controlled and an as-simultaneous-as-possible growth of 5–15 ovarian follicles (*stimulation phase*). When the first three follicles reach a size where a mature oocyte can be expected (about 18 mm in diameter), ovulation is induced, the oocytes are retrieved, those which are mature are tried to be fertilised *in vitro*, and implanted either immediately within the treatment cycle or later after cryopreservation back into the uterus.

Assisted reproduction treatments are complex and challenging, with low average success rates (around 30%) even in the top clinics, and with many factors that, to date, can be hardly kept under full control. Indeed, as hormonal regulatory systems occur within a complex network of endocrinological, neurological and psychological factors [19,16], they are difficult to capture within clinical studies, and model-based approaches might be of great aid in taking these many factors under better control.

Our case study is one of the worldwide classically used downregulation protocols aiming at suppressing the usual hormonal oscillations of the menstrual cycle (as in Figure 1) and preparing the patient to the following stimulation. At the Department of Reproductive Endocrinology of University Hospital Zurich, this protocol currently consists of a sequence of daily administrations of 0.1 mg of Triptorelin (a GnRH analogue). For physiological reasons, downregulation is started within a precise time window of the menstrual cycle, namely between

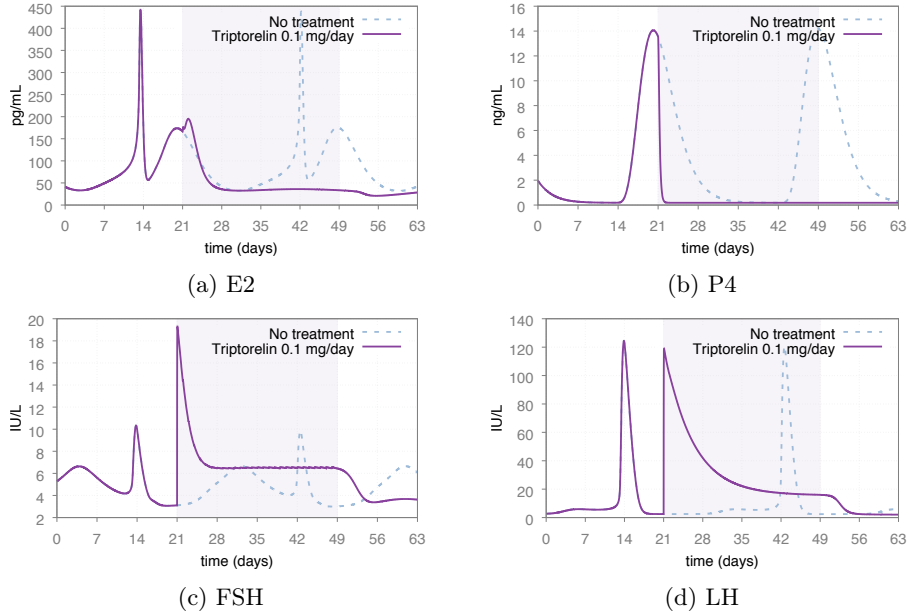


Fig. 1: Expected downregulation effects on some of the hormones playing a role in the human female menstrual cycle. The purple area defines the time window during which Triptorelin is administered.

day 21 and day 25. The treatment might have different duration in order to address different patient reactions. In particular, a downregulation treatment is considered *successful* (effective) for a patient in case the blood concentrations of a given set of hormones and other physiological quantities go below certain thresholds within 9 days from the first drug administration, and stay always below such thresholds for the following 21 days. As a consequence, a downregulation treatment might last up to 30 days.

4 Computing Optimal Personalised Fertility Treatments

In the following, we denote with \mathbb{R} , $\mathbb{R}^{\geq 0}$, \mathbb{R}^+ and \mathbb{N}^+ the sets of, respectively, all real, non-negative real, strictly positive real, and strictly positive natural numbers.

4.1 Formalising the Virtual Physiological Human (VPH) model

VPH models (as GynCycle) are *hybrid systems*. In the typical case (as ours) in which model inputs are drug administrations, VPH models can be abstracted into Discrete Event Systems (DESS) (see, *e.g.*, [23,27]), *i.e.*, continuous-time

input-state-output deterministic causal dynamical systems [38] whose input functions are *discrete event sequences*.

Events for a DES defining a VPH model as GynCycle represent *clinical actions*, and take values from a finite alphabet A (in our case, defining all the possible doses for the treatment drug). We assume that a distinguished element nop exists in A , representing the “null” event (“no action”).

A *discrete event sequence* (defining an input function for our VPH model) is a function $u(t)$ associating to each (continuous) time instant $t \in \mathbb{R}^{\geq 0}$ an event (*i.e.*, a clinical action) in A , and such that $u(t)$ differs from nop only in time points multiple of a given *time quantum* $\tau \in \mathbb{R}^+$. Intuitively, in our setting, a discrete event sequence $u(t)$ defines a series of drug administration actions (including “no action”) occurring at time points multiple of τ .

Given an assignment λ to the VPH model parameters (*i.e.*, a Virtual Phenotype, VP) and a discrete event sequence u , the associated VPH *model trajectory* is a continuous-time function $x(\lambda, t)$ representing the time evolution of the biological quantities defined by the model when fed with u starting from its initial state and with its parameters set to λ . Of course, as the VPH model is causal, its trajectories up to any time point $\bar{t} \in \mathbb{R}^{\geq 0}$ only depend on the *restriction* of the input discrete sequence up to time point \bar{t} [24].

4.2 Modelling treatment invariants and goals

Treatment invariants and goals define conditions that must be, respectively, *always* and *eventually* satisfied by a successful treatment. Invariants and goals for the treatment being sought can be modelled as continuous-time *monitors* embedded within the VPH (DES) model, along the lines of [23,28]. Monitors observe the state of the system and check whether the properties of interest are satisfied.

In particular, given a model trajectory $x(\lambda, t)$ under a given input function $u(t)$, our monitor output is UNDET as long as $x(\lambda, t)$ satisfies the invariants (in other words, the monitor decision is “undetermined” as there is *hope* to extend the current treatment into a successful treatment) and goes to and *stays* at value FAIL as soon as invariants are *violated*. When a UNDET model trajectory satisfies the goal conditions, the monitor output turns to SUCCESS, informing the caller that the input function $u(t)$ defines a *successful* treatment (*i.e.*, a *effective* treatment which *always* satisfies invariants).

The use of continuous-time monitors embedded in the VPH model gives us a flexible way to model both bounded safety and bounded liveness properties (see, *e.g.*, [25,27] for a use of monitors to define safety properties for cyber-physical systems).

In our setting, the properties of interest are the conditions of successful down-regulation treatments (Section 3). In particular, our *invariant* requires that the value of all the biological quantities under observation go and stay beyond their thresholds from the 9-th day after the first drug administration. Our *goal condition* instead requires that values for those quantities stay below their thresholds for 21 consecutive days.




Fig. 2: Treatment monitor.

As a consequence, our monitor output is UNDET in the initial state, and eventually turns either to FAIL or to SUCCESS (see Figure 2). Model output turns to value FAIL as soon as, from the 9-th day after the first drug administration, the value of *any* of the physiological quantities under observation goes beyond its threshold. Model output turns instead from value UNDET to value SUCCESS as soon as all the thresholds are satisfied for 21 consecutive days.

4.3 VPH model simulation

As anticipated in Section 2 the complexity of the (highly non-linear) differential equations typically occurring in actual VPH models hinders the possibility for their symbolic analysis. Numerical *simulation* is often the only means to compute the time evolution of the model quantities of interests (mainly blood hormone concentrations in GynCycle) upon a sequence of clinical actions (administration of one or more drugs with their associated doses). Also, as VPH model equations are often stiff, simulation can be an *expensive* process from a computational point of view.

Our algorithm for optimal personalised treatment computation regards the input VPH model (GynCycle in our case study) as an *executable* black-box model. In order to drive the VPH model simulator, our algorithm assumes that it accepts a set of basic commands (which are available or can be readily implemented within most modern simulators [27]) in order to: (i) set given values (λ , the VP at hand) for the model parameters; (ii) prepare the model to be simulated from its initial state; (iii) inject a clinical action (*i.e.*, administer a drug dose) *and* advance the simulation by time τ (*e.g.*, 1 day); (iv) rollback to a previous simulation state. A *simulation campaign* (see, *e.g.*, [26]) is a sequence of such simulator commands.

4.4 Computing an optimal personalised treatment

In this section we describe the search-based algorithm that we used at the core of our In-Silico Clinical Trial (ISCT).

Given a VP λ of our GynCycle VPH model and a set of possible alternative doses for the drug at hand defining the set A of possible clinical actions (including action “no action”, *i.e.*, dose = 0), our backtracking-based algorithm performs a depth-first search in the space of possible treatments (*i.e.*, discrete sequences of clinical actions in A) lasting at most $h = 25 + 9 + 21 = 55$ days from day 1 of

the patient menstrual cycle (*i.e.*, the latest cycle day, 25, when the treatment can start, plus the latest day, 9, within which the safety conditions must be met, plus the number of consecutive days, 21, in which the safety conditions must be always satisfied in order to declare success).

The search algorithm is implemented in C and drives a GynCycle model simulator (implemented within the Limex Solver [12]) using the commands described in Section 4.3. As the downregulation treatment forbids to administer multiple drug doses in a single day, we can fix the time step τ (time quantum between clinical actions) to 1 day.

The initial state of the GynCycle model represents a patient of VP λ on day 1 of her menstrual cycle. Furthermore, as described in Section 4.2, the model is equipped with a continuous-time monitor which checks at any time whether the current model trajectory satisfies the treatment safety and goal conditions or not.

The algorithm initialises the current treatment to the empty treatment and advances the simulator to cycle day 21 (the earliest day of the first clinical action) with no drug administrations. From this point, at each step it extends the current treatment with a new clinical action $a \in A$ (possibly *nop*, *i.e.*, “no drug”).

Checking safety and goal conditions and backtracking After having appended to the current partial treatment a new action $a \in A$, the action is injected into the simulator, the simulator is advanced by time τ , and the monitor output at the end of the last simulation step is retrieved, thus evaluating the safety and goals treatment conditions (as defined in Section 4.2) after one day from the last clinical action. The simulated (partial) treatment is said *safe* if the monitor output is UNDET, and *unsafe* if the monitor output is FAIL. Whenever the monitor output turns from UNDET to SUCCESS, the current (safe) treatment represents a *successful* treatment. As soon as the monitor returns FAIL, a backtracking is performed, with the previous simulator state being restored (together with the state of the monitor and the monitor output value in the restored state). If no further actions can be executed in the current node of the search tree, a back-track is triggered which restores the simulator (plus monitor) state one level up in the stack.

Optimality constraint Our algorithm does not stop at the first found successful treatment. Indeed, along the lines of [22], it keeps track of the *lightest* successful treatment found so far, *i.e.*, the one envisioning the administration of the *minimum* overall drug amount $D_{\min} \in \mathbb{R}^{\geq 0}$. Initially, D_{\min} is set to $+\infty$.

At each node of the search tree, any action extending the current partial treatment with an administration of a drug dose which would make the overall administered drug amount reaching or exceeding D_{\min} is regarded as not applicable.

Preference ordering among actions In a black-box setting as ours, no inference can be made on the effects of each candidate action without performing a simulator run to actually advance the model and then querying the model monitor output. Hence, approaches to compute, through inference, a dynamic preference order among the candidate actions to be tried during search (as those exploited in, *e.g.*, classical planning, planning for white-box hybrid systems –see [Section 6](#)– or CSP, SAT or local search solvers) cannot be applied.

The only information available to the algorithm without running the simulator is value D_{\min} and the overall cost of the *current* partial treatment (overall amount of drug administered). Hence, given that our algorithm searches for a successful treatment of minimum cost and that any action has a non-negative contribution to the overall treatment cost, it not surprisingly tries the clinical actions on each node in the search tree in ascending order of their associated dose (*i.e.*, cost), hence performing a greedy, optimistic choice.

The optimality constraint (whose threshold is updated each time a new optimal treatment is found), the presence of a bounded horizon h and the fact that our algorithm does not stop at the first successful treatment clearly guarantee that a global optimum is always returned (if any successful treatment exists).

5 Multi-Arm In-Silico Clinical Trial

In this section we outline how we set up and conducted our In-Silico Clinical Trial (ISCT).

5.1 Selection of Virtual Phenotypes (VPs) and exclusion criteria

The GynCycle model defines around 10^4 VPs ([Section 2.1](#)) whose time behaviour (under no treatment) is shown in [Figure 3](#) (light curves). In order to carry out our multi-arm ISCT, we clustered such VPs by merging in the same cluster phenotypes behaving similarly, where similarity has been defined (along the lines of [[39,31](#)]) by means of signal processing metrics. For each obtained cluster, a single representative has been randomly selected.

Not surprisingly, the reference downregulation treatment of [Section 3](#) does not succeed on all possible phenotypes. Hence, we excluded from our ISCT those VPs for which the reference treatment fails. Our *exclusion criterion* directly stems from domain knowledge on downregulation treatments (namely: for these protocols, if the reference treatment, which envisions one full drug dose per day, fails for a patient, a lighter treatment will fail as well).

As a result of the application of clustering and of our exclusion criteria, we obtained a population of 98 VPs representative of the spectrum of all model VPs for which the reference downregulation treatment works (dark curves in [Figure 3](#)). Any two different VPs in our selection differ significantly for the time behaviour of *at least* one of the 33 biological quantities defined by the model (although they might appear similar on others, *e.g.*, those in [Figure 3](#)).

Each of the 98 VPs in our selected population defines a *distinct arm* of our multi-arm ISCT.

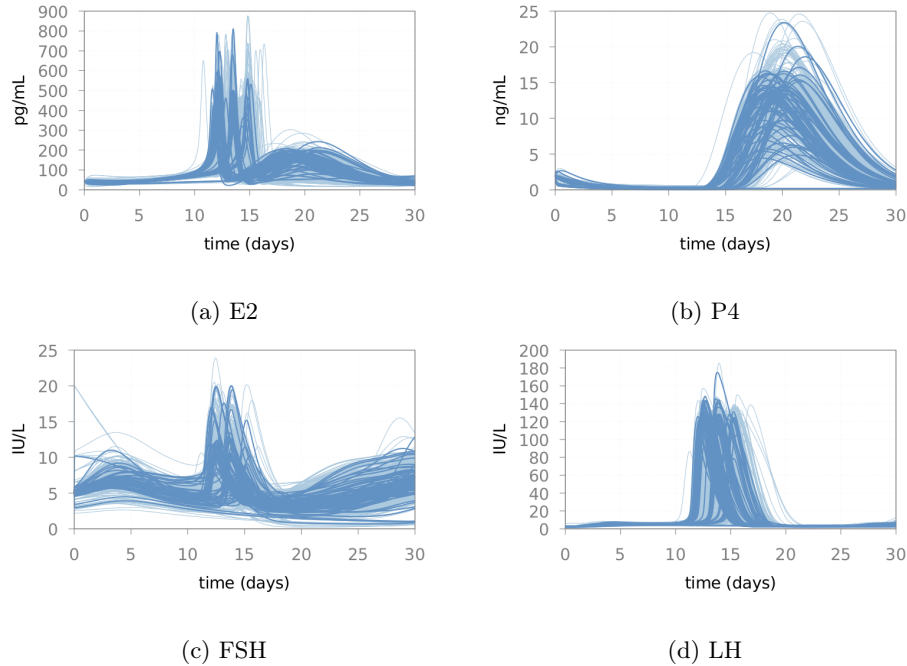


Fig. 3: Time evolutions of all GynCycle VPs for some of the biological quantities defined by the model.

5.2 Multi-arm ISCT run

We ran our 98-arm ISCT using a large High Performance Computing (HPC) infrastructure (the Marconi cluster) kindly provided by the Cineca consortium.

For each VP in the selected population, we ran our algorithm of [Section 4](#) on an independent node of the cluster searching for an optimal (*i.e.*, lightest) downregulation treatment for *that* VP, thus in an *embarrassing parallel fashion*.

5.3 Computational results

[Figure 4a](#) shows the distribution of the computation times of each parallel process implementing a single arm of our multi-arm ISCT.

In particular, the average completion time of the arms of our ISCT is 19.09 h. Although standard deviation is very large (21.31 h), from [Figure 4a](#) it can be seen that the vast majority (around 90%) of ISCT arms terminate in less than 48 hours.

[Figure 4b](#) shows the distribution of the number of computation nodes among our ISCT arms. On average, 918 770 nodes have been expanded by each parallel search process in order to find an optimal treatment for the VP at hand. Also here, standard deviation is very large (1 144 123 nodes).

The high standard deviations of the distributions of computation time and number of expanded nodes among VPs can be explained by observing that our algorithm deterministically orders clinical actions to be selected in each search tree node in ascending order of the administered drug dose (*i.e.*, on their contribution to the overall cost function). In [Section 4.4](#) we have argued that any choice of the actions based on inference of the prospective effect of each action is a not viable option, because of the black-box nature of the Virtual Physiological Human (VPH) model. That said, the left skew of the distributions in [Figures 4a](#) and [4b](#) suggests that our chosen action preference order does behave very well in the *vast majority* of cases.

Overall, [Figures 4a](#) and [4b](#) show the *revolutionary potential* of ISCT: an ISCT over a population of 98 VPs has been conducted in a few days of parallel computation, whilst an equivalent *in vivo* clinical trial would have been not even conceivable. In fact, we would need to recruit in the clinical trial a distinct human patient *for each* phenotype *and for each* of the tested treatment variations. The overall number of human patients to be recruited would have been several millions, thus clearly making the approach infeasible.

5.4 ISCT outcomes

[Figure 5](#) shows the distribution of the total amount of drug employed by our optimal personalised treatments, as found by our algorithm. In particular, our computed personalised treatments use only a *tiny* fraction (on average 16.96%) of the drug quantity employed by the reference treatment of [Section 3](#) (*i.e.*, 0.1 mg on each day for up to 30 days, totalling up to 3 mg and never less than 2.8 mg, since the required thresholds are never reached before day 7).


[Figure 5](#) shows that optimal personalised treatments save, on average, 83.04% of the drug administered in the reference treatment (standard deviation: 2.74%).

However, these extremely favourable results must be read with care. Differently from what happens in clinical practice, our algorithm has *perfect knowledge* about the phenotype of the patient under treatment, hence it can predict *exactly* her reaction to each tentative clinical action (drug administration). As a consequence, both the *dose* and the *timing* of each drug administration can be *optimally* chosen in the personalised treatment.

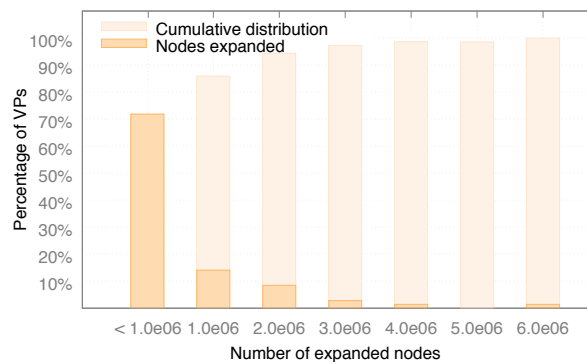
The goal of presenting results as those in [Figure 5](#) is then to show the *potential huge room for improvement* that *in silico* approaches can bring to precision medicine. However, many critical issues are still open in the area of ISCT. Some of them are recapped in [Section 7](#).

6 Related Work

Individualised treatments have the potential value to reduce costs and improve outcomes of standard clinical treatments. In recent years data-driven techniques have been investigated thanks to the availability of big data [\[34\]](#). For example, the knowledge-base approach in [\[13\]](#) has been used to optimise treatment plans



(a)



(b)

Fig. 4: Distributions of computation time (a) and number of expanded search tree nodes (b) among the different ISCT arms.




Fig. 5: Multi-arm ISCT outcomes.

for lung cancer. Unfortunately in presence of scarce experimental data, the above approaches cannot be applied. For example, in our case study hormones blood concentrations are not measured every day, since those measurements are costly and invasive.

Model-based approaches, exploiting Pharmacokinetics (PK), as *e.g.*, [42], are used instead to build virtual phenotypes populations. Such populations are used to optimise and individualise drug doses [18,41]. PK-based models, however, do not define how administered drugs can affect a Virtual Phenotype (VP) (namely, Pharmacodynamics, PD), *i.e.*, possible side-effects due to drug administrations are not taken into account.

In our model-based setting, we have to face with complex Virtual Physiological Human (VPH) models, *e.g.*, HumMod [17], Physiomodel [33] and GynCycle [36] defined through highly non-linear differential equations modelling underlying biological mechanisms (*e.g.*, inhibitory and stimulatory effects).

As outlined in Section 4.1, such VPH models are hybrid systems that can be abstracted into Discrete Event Systems (DESS) (see, *e.g.*, [23,27]) whose inputs are discrete event sequences. To find an optimal treatment means to find an optimal plan in hybrid domains, where the behaviour of the given system is described by both discrete and continuous quantities. In the literature, there are many techniques and tools to model planning problems in hybrid domains. Examples are: PDDL+ [14,40], Satisfiability Modulo Theories (SMT)-based PDDL+ [5] and other PDDL+ extensions [11,10,6,3]. Model checking techniques are also used to find plans. Examples in this direction are [4], which exploits symbolic model checking, UPMurphi [11], which given as input a PDDL+ problem specification computes a universal plan, CGMurphi [9], an explicit model checker used to compute optimal controllers, and [1,32], which define a methodology to compute controllers for non-linear systems.

However, as outlined in Section 4.3, the complexity of the differential equations of VPH models relevant for *in silico* clinical practice makes such models out of reach for symbolic approaches like those mentioned above, and appoints numerical integration as the only viable means to compute (black-box) the model evolutions under a given input function.

In particular, even considering that clinical actions have constant and equal duration (as, *e.g.*, in [20]), no reasoning or inference can be made on action effects in a black-box setting as ours. The automated synthesis of rational decisions and plans in black-box environments, where numerical simulation is the only means to discover the effect of actions, is common in several other application domains of high industrial relevance, like smart grids (see, *e.g.*, [29,30,15]), games (see, *e.g.*, [21]) and real-time manoeuvring of Unmanned Aerial Vehicles (see, *e.g.*, [35]).

7 Conclusions

We presented a case study showing how model-based approaches coupled with intelligent search techniques can be used to support precision medicine by com-

puting, for each given patient, a *personalised treatment* maximising effectiveness while minimising cost as well as likelihood and severity of adverse effects.

By exploiting a complex Virtual Physiological Human (VPH) model of the human female Hypothalamic–Pituitary–Gonadal (HPG) axis, its set of Virtual Phenotypes (VPs), and a backtracking-search algorithm that intelligently drives a simulator for the VPH model, we set up and conducted a multi-arm In-Silico Clinical Trial (ISCT) which computes, for each VP satisfying our inclusion criteria, a personalised variation of a reference pharmaceutical treatment which, preserving effectiveness, minimises the overall amount of drug used.

Our ISCT can be carried out with a few days of parallel computation on a regular High Performance Computing (HPC) infrastructure, while an equivalent (classical) *in vivo* trial would require the recruiting of several millions of human patients in order to test all treatment variations assessed *in silico*, thus resulting just infeasible.

Although our trial shows extremely favourable results (with computed personalised treatments using only a tiny fraction of the drug quantity employed by the reference treatment), there are still open issues in the area of ISCT. In particular, mapping each human patient to her corresponding VP calls for novel and highly inter-disciplinary approaches at the intersection of Artificial Intelligence, Medicine, Mathematics and Biology.

Acknowledgements. This work was partially supported by the Italian Ministry of University and Research (MIUR) under grant “Dipartimenti di eccellenza 2018–2022” of the Department of Computer Science of Sapienza University of Rome and by the EC FP7 project PAEON (Model Driven Computation of Treatments for Infertility Related Endocrinological Diseases, 600773).

References

1. V. Alimiguzhin, F. Mari, I. Melatti, I. Salvo, and E. Tronci. Linearizing discrete time hybrid systems. *IEEE TAC*, 62(10), 2017.
2. M. Bächler, D. Menshykau, Ch. De Geyter, and D. Iber. Species-specific differences in follicular antral sizes result from diffusion-based limitations on the thickness of the granulosa cell layer. *Mol. Hum. Reprod.*, 20(3), 2014.
3. S. Bogomolov, D. Magazzeni, S. Minopoli, and M. Wehrle. PDDL+ planning with hybrid automata: Foundations of translating must behavior. In *Proc. ICAPS 2015*. AAAI, 2015.
4. S. Bogomolov, D. Magazzeni, A. Podelski, and M. Wehrle. Planning as model checking in hybrid domains. In *Proc. AAAI 2014*. AAAI, 2014.
5. D. Bryce, S. Gao, D.J. Musliner, and R.P. Goldman. Smt-based nonlinear pddl+ planning. In *Proc. AAAI 2015*. AAAI, 2015.
6. A.J. Coles and A.I. Coles. Pddl+ planning with events and linear processes. In *Proc. ICAPS 2014*, 2014.
7. L.G.E. Cox, S. Loerakker, M.C.M. Rutten, B.A.J.M. De Mol, and F.N. Van De Vosse. A mathematical model to evaluate control strategies for mechanical circulatory support. *Artif. Organs*, 33(8), 2009.

8. C. Dalla Man, F. Micheletto, D. Lv, M. Breton, B. Kovatchev, and C. Cobelli. The uva/padova type 1 diabetes simulator: New features. *JDST*, 8, 2014.
9. G. Della Penna, B. Intrigila, D. Magazzeni, I. Melatti, and E. Tronci. Cgmurphi: Automatic synthesis of numerical controllers for nonlinear hybrid systems. *Eur. J. Control*, 19(1), 2013.
10. G. Della Penna, B. Intrigila, D. Magazzeni, and F. Mercorio. Planning for autonomous planetary vehicles. In *Proc. ICAS 2010*. IEEE, 2010.
11. G. Della Penna, D. Magazzeni, F. Mercorio, and B. Intrigila. Upmurphi: A tool for universal planning on pddl+ problems. In *Proc. ICAPS 2009*. AAAI, 2009.
12. R. Ehrig, U. Nowak, L. Oeverdieck, and P. Deuffhard. Advanced extrapolation methods for large scale differential algebraic problems. In F.H.-J. Bungartz, editor, *High Performance Scientific and Engineering Computing*, volume 8 of *LNCS*. Springer, 1999.
13. A. Fogliata, F. Belosi, A. Clivio, P. Navarra, G. Nicolini, M. Scorsetti, E. Vanetti, and L. Cozzi. On the pre-clinical validation of a commercial model-based optimisation engine: Application to volumetric modulated arc therapy for patients with lung or prostate cancer. *Radiother. Oncol.*, 113(3):385–391, 2014.
14. M. Fox and D. Long. Modelling mixed discrete-continuous domains for planning. *JAIR*, 27, 2006.
15. B.P. Hayes, I. Melatti, T. Mancini, M. Prodanovic, and E. Tronci. Residential demand management using individualised demand aware price policies. *IEEE Trans. Smart Grid*, 8(3), 2017.
16. M.P. Hengartner, T.H.C. Kruger, K. Geraedts, E. Tronci, T. Mancini, F. Ille, M. Egli, S. Roebnitz, R. Ehrig, L. Saleh, K. Spanaus, C. Schippert, Y. Zhang, and B. Leeners. Negative affect is unrelated to fluctuations in hormone levels across the menstrual cycle: Evidence from a multisite observational study across two successive cycles. *J. Psycho. Res.*, 99, 2017.
17. R. Hester, A. Brown, L. Husband, R. Iliescu, W.A. Pruett, R.L. Summers, and T. Coleman. Hummod: a modeling environment for the simulation of integrative human physiology. *Front. Physiol.*, 2, 2011.
18. P.M. Jeena, W.R. Bishai, J.G. Pasipanodya, and T. Gumbo. In silico children and the glass mouse model: clinical trial simulations to identify and individualize optimal isoniazid doses in children with tuberculosis. *Antimicrob. Agents Chemother.*, 55(2):539–545, 2011.
19. B. Leeners, T.H.C. Kruger, K. Geraedts, E. Tronci, T. Mancini, F. Ille, M. Egli, S. Roebnitz, L. Saleh, K. Spanaus, C. Schippert, Y. Zhang, and M.P. Hengartner. Lack of associations between female hormone levels and visuospatial working memory, divided attention and cognitive bias across two consecutive menstrual cycles. *Front. Behav. Neuro.*, 11, 2017.
20. H.X. Li and B.C. Williams. Generative planning for hybrid systems based on flow tubes. In *Proc. ICAPS 2008*. AAAI, 2008.
21. N. Lipovetzky, M. Ramirez, and H. Geffner. Classical planning with simulators: Results on the atari video games. In *Proc. IJCAI 2015*, volume 15, pages 1610–1616, 2015.
22. T. Mancini, P. Flener, and J. Pearson. Combinatorial problem solving over relational databases: View synthesis through constraint-based local search. In *Proc. SAC 2012*. ACM, 2012.
23. T. Mancini, F. Mari, A. Massini, I. Melatti, F. Merli, and E. Tronci. System level formal verification via model checking driven simulation. In *Proc. CAV 2013*, volume 8044 of *LNCS*. Springer, 2013.

24. T. Mancini, F. Mari, A. Massini, I. Melatti, I. Salvo, and E. Tronci. On minimising the maximum expected verification time. *IPL*, 122, 2017.
25. T. Mancini, F. Mari, A. Massini, I. Melatti, and E. Tronci. Anytime system level verification via random exhaustive hardware in the loop simulation. In *Proc. DSD 2014*. IEEE, 2014.
26. T. Mancini, F. Mari, A. Massini, I. Melatti, and E. Tronci. System level formal verification via distributed multi-core hardware in the loop simulation. In *Proc. PDP 2014*. IEEE, 2014.
27. T. Mancini, F. Mari, A. Massini, I. Melatti, and E. Tronci. SyLVaaS: System level formal verification as a service. In *Proc. PDP 2015*. IEEE, 2015.
28. T. Mancini, F. Mari, A. Massini, I. Melatti, and E. Tronci. Anytime system level verification via parallel random exhaustive hardware in the loop simulation. *Microprocessors and Microsystems*, 41, 2016.
29. T. Mancini, F. Mari, I. Melatti, I. Salvo, E. Tronci, J. Gruber, B. Hayes, M. Prodanovic, and L. Elmegaard. Demand-aware price policy synthesis and verification services for smart grids. In *Proc. SmartGridComm 2014*. IEEE, 2014.
30. T. Mancini, F. Mari, I. Melatti, I. Salvo, E. Tronci, J.K. Gruber, B. Hayes, M. Prodanovic, and L. Elmegaard. User flexibility aware price policy synthesis for smart grids. In *Proc. DSD 2015*. IEEE, 2015.
31. T. Mancini, E. Tronci, I. Salvo, F. Mari, A. Massini, and I. Melatti. Computing biological model parameters by parallel statistical model checking. In *Proc. IWBBIO 2015*, volume 9044 of *LNCS*. Springer, 2015.
32. F. Mari, I. Melatti, I. Salvo, and E. Tronci. Model based synthesis of control software from system level formal specifications. *ACM TOSEM*, 23(1), 2014.
33. M. Matejak and J. Kofranek. Physiomodel-an integrative physiology in modelica. In *Proc. EMBC 2015*. IEEE, 2015.
34. W. Raghupathi and V. Raghupathi. Big data analytics in healthcare: promise and potential. *Health Inf. Sci. Sys.*, 2(1):3, 2014.
35. M. Ramirez, M. Papasimeon, L. Benke, N. Lipovetzky, T. Miller, and A.R. Pearce. Real-time uav maneuvering via automated planning in simulations. In *Proc. IJ-CAI 2017*, pages 5243–5245. AAAI, 2017.
36. S. Roblitz, C. Stotzel, P. Deuffhard, H.M. Jones, D.-O. Azulay, P. van der Graaf, and S.W. Martin. A mathematical model of the human menstrual cycle for the administration of GnRH analogues. *J. Theor. Biol.*, 321, 2013.
37. P.P. Roy and K. Roy. Molecular docking and qsar studies of aromatase inhibitor androstenedione derivatives. *J. Pharm. Pharmacol.*, 62(12):1717–1728, 2010.
38. E.D. Sontag. *Mathematical Control Theory: Deterministic Finite Dimensional Systems (2nd Ed.)*. Springer, 1998.
39. E. Tronci, T. Mancini, I. Salvo, S. Sinisi, F. Mari, I. Melatti, A. Massini, F. Davi, T. Dierkes, R. Ehrig, S. Roblitz, B. Leeners, T. H. C. Kruger, M. Egli, and F. Ille. Patient-specific models from inter-patient biological models and clinical records. In *Proc. FMCAD 2014*. IEEE, 2014.
40. M. Vallati, D. Magazzeni, B. De Schutter, L. Chrapa, and T.L. McCluskey. Efficient macroscopic urban traffic models for reducing congestion: A pddl+ planning approach. In *Proc. AAAI 2016*, pages 3188–3194, 2016.
41. S.C. van Dijkman, S.G. Wicha, M. Danhof, and O.E. Della Pasqua. Individualized dosing algorithms and therapeutic monitoring for antiepileptic drugs. *Clin. Pharmacol. Ther.*, 103(4):663–673, 2018.
42. S. Willmann, J. Lippert, M. Sevestre, J. Solodenko, F. Fois, and W. Schmitt. Pk-sim[®]: a physiologically based pharmacokinetic ‘whole-body’ model. *BIOSILICO*, 1(4):121–124, 2003.